

Applicant : Donald F. Hooper
Serial No. : 09/608,354
Filed : June 29, 2000
Page : 7 of 12

Attorney's Docket No.: 10559-222001 / P8715
Intel Corporation

REMARKS

Applicant amended independent claim 23 to clarify that the data structure used for the longest prefix match program includes a first table at the root that whose entries include links to a first set of trie tables, and a second table at the root whose entries include route pointers, some of which pointing to route entries, and trie pointers, some of which point to a second set of trie tables with additional index entries containing route and trie pointers.

Independent claim 23 was also amended for greater clarity. Support for the clarification is found, for example, at FIGS. 3 and 4, and page 5, line 19 to page 7, line 12. Independent claim 28 was similarly amended. In addition, claims 24, 26, 29-30 and 32-35 were all amended to make the language recited in those claims consistent with the language recited in the amended independent claims. Furthermore, applicant added new claims 36 and 37 reciting features similar to those recited in original claims 21 and 22 presented in the originally filed application. After these amendments claims 23-37 are pending. Claims 23 and 28 are independent.

The Examiner rejected claims 23-35 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 6,192,051 to Lipman et al., in view of U.S. Patent No. 6,493,347 to Sindhu et al. and U.S. Patent No. 6,691,218 to Brown.

Applicant's amended independent claim 23 recites a router that includes "a data structure stored in the memory, the data structure including information used by the longest prefix match program and including: at least one first table at a root, the at least one first table comprising entries indexed according to a first portion of an Internet Protocol (IP) address span and maintaining links to a first set of nodes containing a first set of trie tables corresponding to less significant portions of the IP address span, the entries of the at least one first table not including links to route entries, each table of the first set of trie tables configured to be indexed according to address bits in the corresponding IP address span of that trie table to locate an indexed trie entry, the indexed trie entry including a corresponding route pointer and a corresponding trie pointer; and at least one second table at the root, the at least one second table comprising entries indexed according to a second portion of the IP address span, each entry of the at least one second root table including a corresponding route pointer and a corresponding trie pointer, at

Applicant : Donald F. Hooper
Serial No. : 09/608,354
Filed : June 29, 2000
Page : 8 of 12

Attorney's Docket No.: 10559-222001 / P8715
Intel Corporation

least some of the entries having corresponding trie pointers pointing to entries in a second set of trie tables at a second set of nodes."

Use of one root table providing links to trie tables, but not route entry pointers, as well as another root table having route entry pointers and links to other trie tables, enables a quick and efficient determination of the longest prefix match for a particular IP address. Advantageously, applicant's data structure greatly reduces table update time.

In contrast, none of the prior art references cited by the examiner discloses applicant's data structure, as recited in independent claim 23.

Lipman describes a network routing apparatus that employs a multi-level tree data structures in a centralized routing table and in distributed forwarding tables (Abstract). Lipman's routing apparatus includes a system controller 12 that performs routing management operations, and thus maintains a large routing table which enables router 10 to make decisions on how packets received at routers 10 should be forwarded (FIG. 1, and col. 7, lines 17-27). With respect to the routing table maintained on system controller 12, Lipman further describes:

The routing table contains a plurality of routing entries, segregated into groups as follows: level-1 routing entries 138 which have subnet mask lengths of 16 or less; level-2 routing entries 140 having subnet mask lengths of between 17 and 24; and level-3 routing entries 142 having subnet mask lengths of between 25 and 32.

Each routing entry 138, 140 and 142 includes a key and a "next hop" address. The key includes a destination IP address value and a subnet mask value (shown together in the form X.X.X.X/yy at the top of each routing entry). The next hop address is an IP address (shown in the form X.X.X.X at the bottom of each routing entry) which is the address of a network to which frames having destination addresses matching the key should be forwarded.

...

Each pointer storage location in the level-1 tree 144 contains either a pointer or a null value indicating that the location is empty. In addition, each location contains additional bits that describe the pointer type when a pointer is present. Each pointer points to a corresponding routing entry and each routing entry may be part of a linked list of routing entries as described above. A pointer may point to a routing entry whose key address value is equal to the index of the location at which the pointer is stored. Such pointers are termed "real leaf" pointers, and this status is indicated by a bit called the "real leaf" bit being set to one. A pointer may instead point to a routing entry whose key address value is not equal to the index of the pointer, but is the best match for that index based on the subnet mask.

Applicant : Donald F. Hooper
Serial No. : 09/608,354
Filed : June 29, 2000
Page : 9 of 12

Attorney's Docket No.: 10559-222001 / P8715
Intel Corporation

These pointers are termed "fill leaf" pointers, as indicated by the "real leaf" bit being set to zero. In FIG. 7, the path between a fill pointer and the associated routing entry is shown using a dotted line. (Lipman, FIG. 7 and col. 10, line 32, to col. 11, line 17)

Lipman routing table includes a single root table having entries that include both route entries and pointers to entries in other tree levels. Lipman does not disclose or suggest applicant's recited data structure which includes "at least one first table at a root, the at least one first table comprising entries indexed according to a first portion of an Internet Protocol (IP) address span and maintaining links to a first set of nodes containing a first set of trie tables corresponding to less significant portions of the IP address span, the entries of the at least one first table not including links to route entries," and/or "at least one second table at the root, the at least one second root table comprising entries indexed according to a second portion of the IP address span," as required by applicant's independent claim 23.

Sindhu describes a router for switching data packets from a source to a destination in a network, in which distributed memory banks are used to store uniform portions of received data packets (Abstract). Although Sindhu discloses a controller 106 that performs a trie-based search based on key information read from the data packets to identify the output ports through which the data packets should be routed (col. 6, lines 9-30), Sindhu is silent on how its trie structures are configured or arranged, or what data are stored in the trie structures, and does not disclose or suggest a data structure that includes "at least one first table at a root, the at least one first table comprising entries indexed according to a first portion of an Internet Protocol (IP) address span and maintaining links to a first set of nodes containing a first set of trie tables corresponding to less significant portions of the IP address span, the entries of the at least one first table not including links to route entries," and/or "at least one second table at the root, the at least one second root table comprising entries indexed according to a second portion of the IP address span," as required by applicant's independent claim 23.

Brown describes a lookup table that provides a longest prefix match for a key provided in a search request (Abstract). Particularly, Brown's lookup table 100 includes binary tree

Applicant : Donald F. Hooper
Serial No. : 09/608,354
Filed : June 29, 2000
Page : 10 of 12

Attorney's Docket No.: 10559-222001 / P8715
Intel Corporation

representations stored in mappers 106a-c (FIG. 2A, col. 5, lines 46-48). Brown further describes:

The value of any node in the tree can be determined by tracing a path from the root 114. Each node in the binary tree is shown with two children, a right child and a left child. The right child is chosen if the parent node is '1'. The left child is chosen if the parent node is '0'. Tracing the path from the root 114 to node 116, r1 is stored as the route index 102 in the L1 mapper 106a for all keys with MSBs set to '010'. Tracing the path from the root node 114 to s0 node 130.sup.4, s0 is stored in the L1 mapper 106a for all keys with MSBs set to '00011'.

...
Each location in the L1 mapper 106a stores a route index 102 assigned to the level-5 node 300.sup.1 -300.sup.32 directly or through a parent of the level-5 node 300.sup.1-32 or an address pointer to the next mapper 106b-c. Mapper level_3112c includes two host nodes h0 at node 138 and h1 at node 140 at the bottom level of the 32-level binary tree. A search for a host node requires a search of all bits of the key 104. As shown in FIG. 2A the route index 102 for h0 is stored in location 146.sup.4 in L3_mapper 106c. Unlike the L1 mapper 106a, the L2 mapper 106b and L3 mapper 106c are not directly mapped. (col. 6, lines 21-62)

Brown's lookup table includes a root data structure that includes entries having route entries and pointers to another level of the lookup table structure (namely, mappers 106b-c). Brown's lookup table, however, does not use trie tables. Further, although Brown describes that several lookup tables (e.g., 200a-b) may be used to search for the value corresponding to a search key, the multiple lookup tables searched are structured similarly to the lookup table 100 described above. Accordingly, Brown does not disclose or suggest applicant's recited data structure that includes at least the features of "at least one first table at a root, the at least one first table comprising entries indexed according to a first portion of an Internet Protocol (IP) address span and maintaining links to a first set of nodes containing a first set of trie tables corresponding to less significant portions of the IP address span, the entries of the at least one first table not including links to route entries," and/or "at least one second table at the root, the at least one second table comprising entries indexed according to a second portion of the IP address span, each entry of the at least one second table including a corresponding route pointer and a corresponding trie pointer, at least some of the entries having corresponding trie pointers pointing to entries in a second set of trie tables at a second set of nodes."

Applicant : Donald F. Hooper
Serial No. : 09/608,354
Filed : June 29, 2000
Page : 11 of 12

Attorney's Docket No.: 10559-222001 / P8715
Intel Corporation

Since none of Lipman, Sindhu or Brown discloses or suggests, alone or in combination, at least the feature of "a data structure stored in the memory, the data structure including information used by the longest prefix match program and including: at least one first table at a root, the at least one first table comprising entries indexed according to a first portion of an Internet Protocol (IP) address span and maintaining links to one or more first set of nodes containing a first set of trie tables corresponding to less significant portions of the IP address span, the entries of the at least one first table not including links to route entries," and/or "at least one second table at the root, the at least one second table comprising entries indexed according to a second portion of the IP address span, each entry of the at least one second table including a corresponding route pointer and a corresponding trie pointer, at least some of the entries having corresponding trie pointers pointing to entries in a second set of trie tables at a second set of nodes," applicant's independent claim 23 is therefore patentable over the cited art. Claims 24-27 and new claims 36-37 depend from independent claim 23 and are therefore patentable for at least the same reasons as independent claim 23.

Independent claim 28 recites "a data structure stored in the memory, the data structure including information used by the longest prefix match program and including: at least one first table at a root, the at least one first table comprising entries indexed according to a first portion of an Internet Protocol (IP) address span and maintaining links to one or more first set of nodes containing a first set of trie tables corresponding to less significant portions of the IP address span, the entries of the at least one first table not including links to route entries," and/or "at least one second table at the root, the at least one second table comprising entries indexed according to a second portion of the IP address span, each entry of the at least one second table including a corresponding route pointer and a corresponding trie pointer, at least some of the entries having corresponding trie pointers pointing to entries in a second set of trie tables at a second set of nodes." For similar reasons as those provided with respect to independent claim 23, at least this feature is not disclosed by the cited art. Claim 29-35 depend from independent claim 28 and are therefore patentable for at least the same reasons as independent claim 28.

Applicant : Donald F. Hooper
Serial No. : 09/608,354
Filed : June 29, 2000
Page : 12 of 12

Attorney's Docket No.: 10559-222001 / P8715
Intel Corporation

It is believed that all the rejections and/or objections raised by the examiner have been addressed.

All of the dependent claims are patentable for at least the reasons for which the claims on which they depend are patentable.

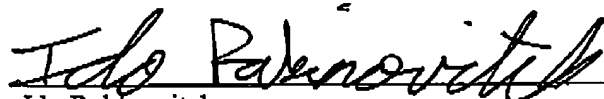
Canceled claims, if any, have been canceled without prejudice or disclaimer.

Any circumstance in which the applicant has (a) addressed certain comments of the examiner does not mean that the applicant concedes other comments of the examiner, (b) made arguments for the patentability of some claims does not mean that there are not other good reasons for patentability of those claims and other claims, or (c) amended or canceled a claim does not mean that the applicant concedes any of the examiner's positions with respect to that claim or other claims.

Please apply any charges to deposit account 06-1050, referencing attorney docket 10559-222001.

Respectfully submitted,

Date: Feb. 21, 2006


Ido Rabinovitch
Attorney for Intel Corporation
Reg. No. L0080

PTO Customer No. 20985
Fish & Richardson P.C.
Telephone: (617) 542-5070
Facsimile: (617) 542-8906

21270724.doc